Understanding EA Dynamics via Population Fitness Distributions

Elena Popovici and Kenneth De Jong

Department of Computer Science
George Mason University
Fairfax, VA 22030

1 Motivation and Methodology

This paper introduces a new tool to be used in conjunction with existing ones for a more comprehensive understanding of the behavior of evolutionary algorithms. Several research groups including [1,3,4] have shown how deeper insights into EA behavior can be obtained by focusing on the changes to the entire population fitness distribution rather than just "best-so-far" curves. But characterizing how repeated applications of selection and reproduction modify this distribution over time proved to be very difficult to achieve analytically and was done successfully for only a few very specialized EAs and/or very simple fitness landscapes.

Our approach is to study empirically derived fitness distributions, both qualitatively and quantitatively, believing that they have the potential for providing interesting and useful insights into the behavior of EAs. The methodology is quite general and can be applied to any EA and to any fitness landscape.

We instrument an EA to provide snapshots of the population fitness data at designated points during the evolutionary process, and display them in a histogram-like manner. While visualizing population fitness distributions can be quite insightful, we go a step further and perform a quantitative analysis of the observed fitness distributions, by estimating how likely the empirically generated distributions reflect an underlying standard distribution (e.g. normal, exponential, etc.). The statistical technique used was the generation of "Q-Q" plots from which we computed an R^2 value representing the likelihood that an empirically generated distribution is due to a certain theoretical distribution [2].

2 Experiments and Results

Our initial experiments were designed to provide some insight into the following questions: 1) how dependent are these observations on the particular type of EA being used, 2) how dependent are the observed population fitness distributions on a particular fitness landscape, and 3) can useful insights be gained by taking snapshots within a generation to study the effects of selection and reproduction on population fitness distributions?

As earlier work had focused on the population fitness distributions of standard generational GAs, for contrast we focused on an EA using a binary representation, standard crossover and mutation, but $\mu + \lambda$ population dynamics.

E. Cantú-Paz et al. (Eds.): GECCO 2003, LNCS 2724, pp. 1604–1605, 2003.

We performed experiments on two problems: 1) the F1 function of the De Jong test suite, used with dimensionality 4 on $[-5.12, 5.12]$ and 2) evolving rules for two-dimensional cellular automata that generate some predefined pattern.

Our experiments showed that the $\mu + \lambda$ EA, due to its strong truncation selection pressure, dramatically distorts the shape of the fitness distributions and steadily decreases their mean and variance. By contrast, a GA on the same problems reaches a steady state rather quickly.

The various operators used have different effects on the fitness distribution as well. There is little difference in the distortions produced by crossover and mutation in the early generations. But, as the population converges, the difference in distortions becomes more apparent (crossover fails to match any of the theoretical distributions while mutation still fits some of them quite well).

Experiments with the CA domain show the dependency of fitness distributions on the landscape. However, the interesting thing is that although the population fitness distributions for the CA problem are quite different in shape, the distortion effects due to crossover and mutation are much the same as we saw on F1. In the early generations they have pretty much the same effect, but increasingly differ in their effects as evolution proceeds.

3 Conclusions

In summary we are optimistic that the methodology presented here will prove to be a useful addition to the current set of tools for analyzing the behavior of EAs. Even the simple experiments presented here have yielded useful insights, including the fact that population fitness distributions are *seldom* observed to be normally distributed, that the shapes of these distributions are heavily dependent on both the fitness landscape and the EA selection pressure, and that the differences between the population fitness distribution distortions due to crossover and mutation are only significant in the later stages of the evolutionary process. In addition, using this methodology, one can actually verify the validity of various assumptions about fitness distributions that theoretical work in this area makes in order to keep the mathematics tractable.

References

1. Lee Altenberg. The Schema Theorem and Price's Theorem. In L. Darrell Whitley and Michael D. Vose, editors, *Foundations of Genetic Algorithms 3*, pages 23–49, Estes Park, Colorado, USA, 1995. Morgan Kaufmann.
2. R. Jain. *The Art of Computer Systems Performance Analysis.* John Wiley and Sons, Inc., New York, 1991.
3. H. Mühlenbein and D. Schlierkamp-Voosen. Predictive models for the breeder genetic algorithm. *Evolutionary Computation*, 1(1):25–49, 1993.
4. J. Shapiro, A. Prügel-Bennett, and M. Rattray. A statistical mechanical formulation of the dynamics of genetic algorithms. In Terence C. Fogarty, editor, *Evolutionary Computing, AISB Workshop*, volume 993 of *Lecture Notes in Computer Science*. Springer, 1994.